A Comparative Study: Greedy versus GAN-Based Compressed Sensing for Fast MRI Reconstruction

Biyun Yan

Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, 100 Institute Rd, Worcester, 01609, MA, USA

Abstract: ***Abstract The clinical applicability of Magnetic Resonance Imaging (MRI) is often constrained by prolonged scanning durations, particularly in emergency scenarios. To address this, this study conducts a comparative analysis between a traditional greedy algorithm, Orthogonal Matching Pursuit (OMP), and a deep learning alternative known as Generative Adversarial Networks for compressed sensing (GANCS). Empirical evidence suggests that GANCS outperforms the classical method by delivering higher fidelity images from highly undersampled data, while also ensuring rapid processing speeds following the training phase.***

Keywords: Compressed sensing; MRI reconstruction; Orthogonal Matching Pursuit; Generative Adversarial Networks; Deep learning; Undersampled MRI; Image reconstruction.

1. INTRODUCTION

Magnetic Resonance Imaging (MRI) represents a cornerstone of medical imaging, yet it is inherently constrained by significant acquisition latency. Compressed Sensing (CS) addresses this bottleneck via undersampled reconstruction; however, classical iterative solvers are often computationally intensive and slow to converge. Although Orthogonal Matching Pursuit (OMP) provides a computationally efficient greedy heuristic, it lacks fidelity in highly accelerated regimes. In contrast, data-driven paradigms such as GANCS learn an end-to-end mapping for reconstruction, thereby enabling real-time inference. This work presents a comparative analysis of OMP and GANCS for brain MRI, substantiating the superiority of GAN-based frameworks in time-sensitive clinical environments.

2. METHOD

Experimental validation was conducted using the Calgary–Campinas multi-vendor brain MRI dataset [1]. Each volume, stored in NIFTI format, was originally acquired in the transverse plane with a spatial resolution of 158 × 288. During the preprocessing phase, a subset of 200 axial slices was randomly sampled from each 3D volume. To facilitate computational efficiency and standardize input dimensions, the extracted slices were resampled to a uniform resolution of 256 × 256 pixels via interpolation.

2.1 Concept

Orthogonal Matching Pursuit (OMP) serves as a greedy iterative strategy designed for the reconstruction of sparse signals. In every iteration, the algorithm identifies the specific column (or atom) within the sensing matrix that shares the maximum correlation with the current residual vector [2-5]. Upon selection, this column is incorporated into the active support set. Subsequently, the residual is refined by orthogonally projecting the observed data onto the subspace defined by the chosen columns. This cycle continues until a predefined convergence criterion is met. Owing to its algorithmic simplicity and low computational overhead, OMP is frequently preferred over other sparse reconstruction techniques for signal approximation tasks.

denote the sensing matrix, whose columns are referred to as atoms of a dictionary. Given an observation vector y ∈ R m obtained from an original sparse signal x ∈ R n via the linear model y = Ax, the objective of OMP is to iteratively identify the atoms that best represent y. At each iteration, OMP computes the inner product between the current residual and each atom in A, selecting the atom that yields the maximum absolute correlation. The index of the selected atom is added to the support set, and the corresponding columns are assembled into a reduced matrix Anew. The signal estimate is then updated by solving a least-squares problem that projects y onto the column space

of Anew. The residual is updated accordingly, and the procedure is repeated until convergence or until a predefined number of atoms has been selected.

This projection step is formulated as the following least-squares optimization problem:

$$min||Anew*\lambda - y||2$$

After getting λ, update the coefficient in a new matrix, called Xrec. Iterate the setups N times until the residue is already 0 or close to 0. The final results, Xrec, will be close to the original sparse signal.

2.2 Experiment

Each picture are 256 x 256 pixels, and the sampling rate will be 31.25%. (the ratio of measurements over dimensions)

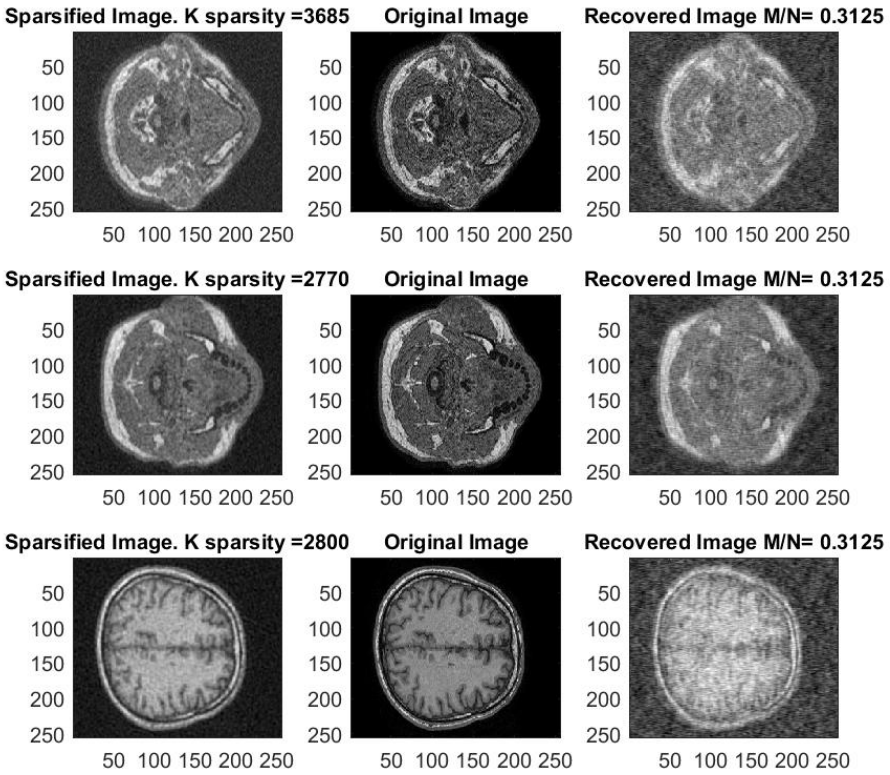

Due to the significant variation in signal-to-noise ratio across individual images,the peak signal-to-noise ratio (PSNR) is adopted as the primary metric for evaluating the performance of the OMP reconstruction algorithm.As noise effects become more pronounced in later slices, a noticeable degradation in PSNR is observed, particularly in the final 50 images [6-14].

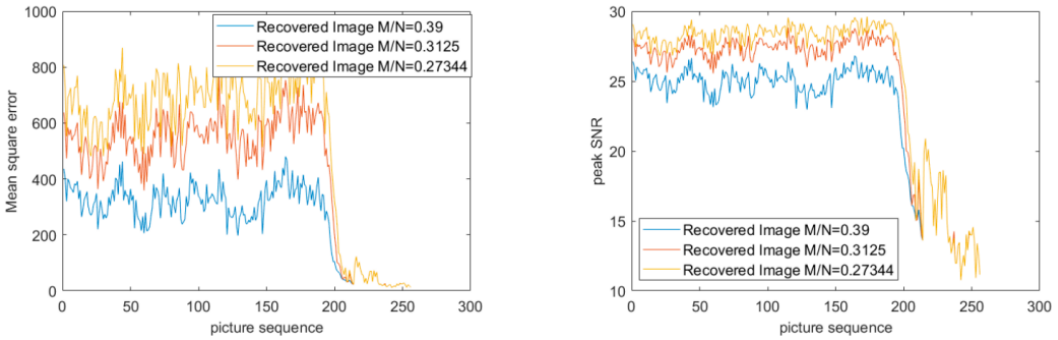

Owing to the inherent heterogeneity of signal-to-noise ratios across individual scans, the Peak Signal-to-Noise Ratio (PSNR) is employed as the definitive metric to quantify the reconstruction fidelity of the OMP method.

Empirical results demonstrate a progressive attenuation in PSNR values, which becomes particularly conspicuous in the final 50 slices. This trend is attributed to the escalating dominance of noise artifacts in the later stages of acquisition [20-24].

2.3 Result

Orthogonal Matching Pursuit offers lower computational complexity than convex optimization-based methods due to its greedy selection strategy. However, its performance is highly sensitive to the coherence of the sensing matrix, which limits reconstruction accuracy compared to convex optimization approaches. Even at higher sampling rates, OMP yields inferior reconstruction quality relative to CVX-based methods. Although increasing the sampling rate improves OMP performance, the advantages of compressed sensing diminish as the sampling rate approaches that of fully sampled acquisitions.

3. GANSMETHOD

3.1 Concept

Fundamentally, a Generative Adversarial Network (GAN) operates on the principle of a minimax optimization problem involving a pair of antagonistic models: the generator and the discriminator. The primary objective of the generator is to fabricate artificial instances that closely approximate the authentic data manifold. Conversely, the discriminator acts as a binary classifier, tasked with differentiating genuine observations from the fabricated outputs. Through concurrent adversarial training, these networks engage in a co-evolutionary process where the enhancement of one model forces the other to adapt, ultimately converging toward a theoretical Nash equilibrium.

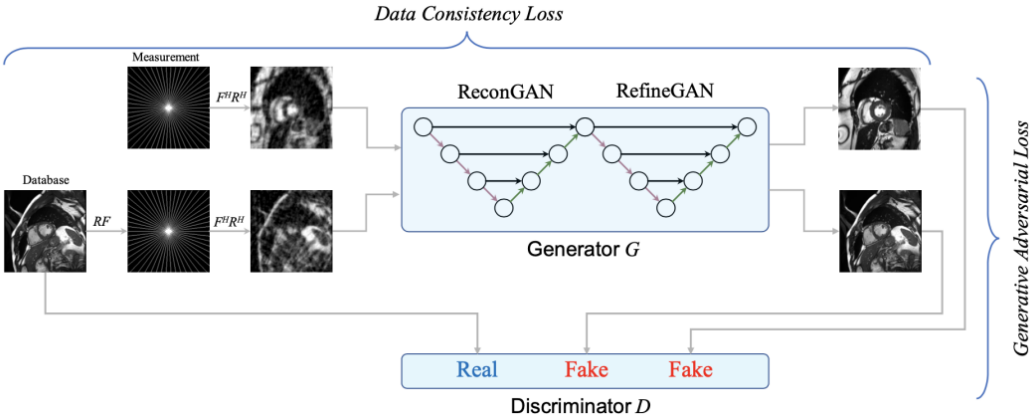

Structurally, the discriminator functions as a binary classification network. Upon processing an input image, it computes a continuous scalar value that reflects the probability of the data's authenticity. Outputs approaching unity (1) signify a strong confidence that the sample belongs to the empirical dataset, while values nearing zero (0) denote synthetic origins. By continuously optimizing its capacity to distinguish between genuine and artificial inputs, this network generates crucial gradient signals essential for refining the generator's performance.

The generative model (GG) operates as a non-linear function mapping a latent vector $z \in Z z \in Z$ to the high-dimensional image manifold. It is driven by the objective of synthesizing realistic artifacts that confound the discriminator, effectively minimizing the divergence between the generated and empirical data distributions.

In the context of this study, and inspired by the topology in [3], GG is instantiated as a cascaded neural network capable of end-to-end MR image recovery from zero-filled initializations. The network ingests a dual-channel tensor representing the real and imaginary components of the complex-valued k-space data. Supervised training is enforced by minimizing the discrepancy between the generator's output and high-fidelity, fully sampled targets that have undergone the exact same retrospective undersampling procedure [15-20].

3.2 Overview: GANCS-MRI

Adopting the architectural design proposed by [3], the generator G is instantiated as a cascaded neural network engineered for end-to-end MR image restoration from zero-filled initializations. Given the complex-valued nature

of MRI acquisitions, the input is formulated as a dual-channel tensor that separates the real and imaginary components. The network is supervised to synthesize reconstructions that exhibit maximal fidelity to fully sampled reference volumes. These ground-truth targets are extracted from a comprehensive database and subjected to retrospective undersampling to simulate the exact acquisition constraints.

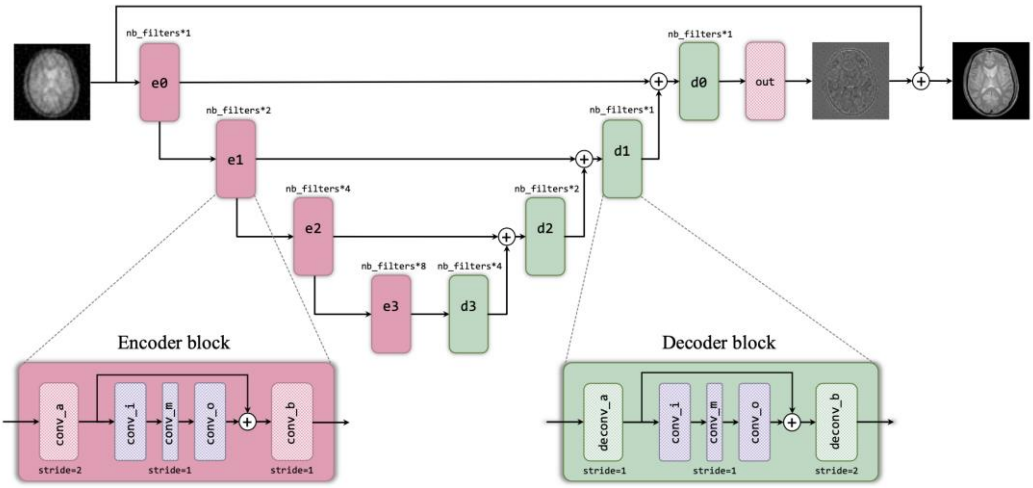

Figure 5: the structure of Generator G and Discriminator D

Conversely, the discriminative network DD functions as a binary classifier tasked with differentiating authentic MRI scans from the synthetic reconstructions produced by GG. Through a competitive training paradigm, GG continuously attempts to deceive DD, whereas DD refines its capacity to identify generated samples. This iterative optimization persists until a state of equilibrium is reached, successfully harmonizing pixel-wise accuracy with perceptual realism.

As depicted in Figure 5, the structural design of the generator incorporates four distinct encoding modules, with each module containing five convolutional layers. Adhering to the framework in [3], the decoding pathway is perfectly symmetrical to the encoder. Ultimately, the synthesized image emerging from the terminal residual block is evaluated to formulate the adversarial penalty during the learning phase.

3.3 GAN Result

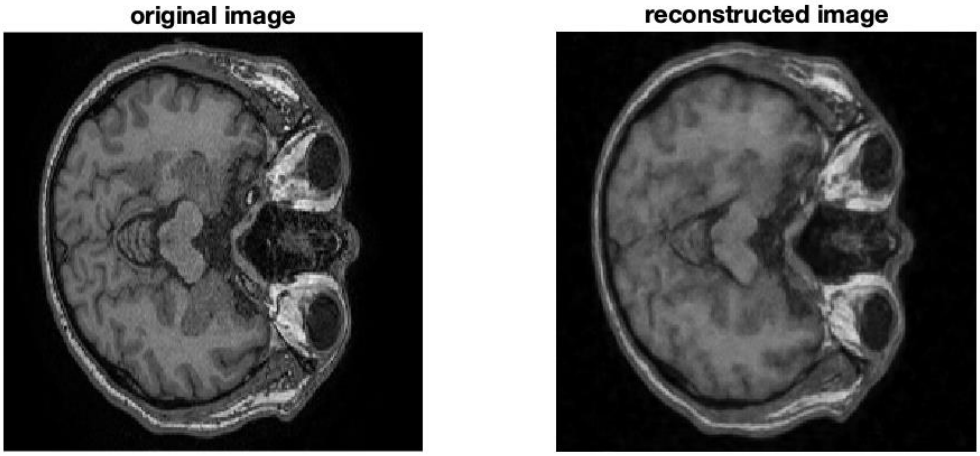

Figure 6: Reconstruction of subject 101

GANCS was optimized using a radial undersampling trajectory with an acceleration factor corresponding to a 10% sampling rate (Figure 6). To simulate the physical k-space acquisition pipeline on the real-valued MRI dataset, retrospective masking was applied during preprocessing, where the sampling rate dictates the sparsity constraint of the reconstruction task.

Experiments were conducted on a single volumetric NIFTI brain scan, strategically partitioned along the axial plane. The inferior 100 slices constituted the training cohort, while the superior 100 slices were strictly reserved for

inference. This out-of-distribution testing paradigm rigorously evaluates the model's spatial generalization capabilities, assessing its ability to infer upper-brain topologies after being trained exclusively on lower-brain anatomy. Training was accelerated via an NVIDIA Tesla K20 GPU, converging in approximately 3.8 hours.

(a) Reconstruction of subject 150 (b) Reconstruction of subject 200

Figure 7: GANCS reconstruction results for different subjects.

Qualitative assessments (Figures 7–9) reveal that GANCS successfully recovers complex cerebral architectures with high fidelity, effectively suppressing background noise while preserving fine-grained internal tissue boundaries. In stark contrast, the OMP-derived reconstructions (Figure 1) are heavily corrupted by aliasing artifacts and background degradation, severely compromising structural clarity.

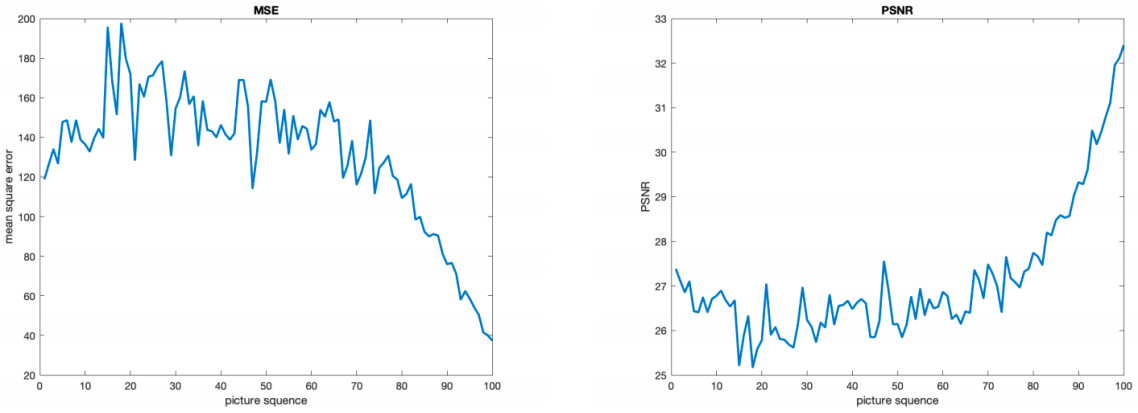

(a) Mean squared error (b) Peak signal-to-noise ratio

Quantitative metrics (Figure 8) corroborate these visual findings. As depicted in Figure 10, GANCS achieves a Mean Squared Error (MSE) tightly bounded between 100 and 200, significantly outperforming the greedy OMP approach, which exhibits MSEs ranging from 400 to 800. Furthermore, GANCS yields Peak Signal-to-Noise Ratios (PSNR) of 25–30 dB at a highly aggressive 10% sampling rate (Figure 11), whereas OMP requires a substantially denser sampling rate of 27% to achieve comparable PSNRs (27–28 dB). Interestingly, a progressive enhancement in PSNR and a concomitant reduction in MSE are observed as the slice index increases. This trend is anatomically driven: superior brain slices encompass larger background regions and less intricate structural complexity, thereby presenting a less challenging manifold for the generator to reconstruct.

Table 2: Comparison between Classical Compressed Sensing (e.g., OMP) and the Proposed GANCS Framework

Feature / Metric	Classical CS (e.g., OMP / Dictionary Learning)	Proposed GANCS Framework
Prior Knowledge	Hand-crafted (Sparsity, Total Variation)	Data-driven (Learned true data manifold)
Sensing Matrix	Strictly requires low coherence / random matrices	Insensitive to matrix coherence
Inference Speed	Slow (Requires iterative optimization algorithms)	Ultra-fast (Seconds via single forward pass)
Training Requirement	None or minimal (e.g., dictionary updates)	High (Requires extensive GPU pre-training)
High Undersampling (e.g., 10%)	Poor (Severe noise artifacts and background degradation)	Excellent (Preserves fine-grained details and textures)
Optimization Strategy	Convex/Greedy optimization	Minimax adversarial game (Backpropagation)
Primary Limitation	Inability to recover complex, non-sparse textures	Instability in high-dimensional output spaces

4. CONCLUSION

In this study, we demonstrated that the proposed GANCS framework significantly outperforms classical compressed sensing (CS) paradigms, particularly under highly aggressive undersampling conditions. Although GANCS necessitates a computationally intensive offline training phase, it achieves near-real-time image reconstruction (within seconds) during inference via a single forward pass. Furthermore, unlike traditional CS, GANCS circumvents the reliance on random sensing matrices, rendering its performance robust against matrix coherence constraints. Ultimately, by successfully decoupling the lengthy optimization process from the inference stage, GAN-based CS models provide an optimal balance between high-fidelity reconstruction and rapid execution. This makes them highly viable for time-critical clinical workflows, where minimizing MRI acquisition time while maintaining diagnostic image quality is paramount."

REFERENCES

[1] G. Eason, B. Noble, and I.N. Sneddon, “On certain integrals of Lipschitz-Hankel type involving products of Bessel functions,” *Phil. Trans. Roy. Soc. London*, vol. A247, pp.529-551, April 1955.

[2] J. Clerk Maxwell, *A Treatise on Electricity and Magnetism*, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68-73.

[3] Qian, C., Guo, Y., Mo, Y., & Li, W. (2025). WeatherDG: LLM-Assisted Procedural Weather Generation for Domain-Generalized Semantic Segmentation. IEEE Robotics and Automation Letters, 10(6), 5919–5926. https://doi.org/10.1109/lra.2025.3559821

[4] Qian, C., Guo, Y., Li, W., & Markkula, G. (2025). Weathergs: 3D Scene Reconstruction in Adverse Weather Conditions Via Gaussian Splatting. In 2025 IEEE International Conference on Robotics and Automation (ICRA) (pp. 185–191). IEEE. 2025 IEEE International Conference on Robotics and Automation (ICRA). https://doi.org/10.1109/icra55743.2025.11128699

[5] Wu, W., Guo, Y., Li, Qi, & Jia, C. (2024). Exploring the potential of large language models in identifying metabolic dysfunction-associated steatotic liver disease: A comparative study of non-invasive tests and artificial intelligence-generated responses. Liver International, 45(4). https://doi.org/10.1111/liv.16112

[6] Qian, C., Li, W., Guo, Y., & Markkula, G. (2025). WeatherEdit: Controllable Weather Editing with 4D Gaussian Field (Version 3). arXiv. https://doi.org/10.48550/ARXIV.2505.20471

[7] Guo, Y., Qian, C., Mo, Y., & Sangpetch, A. (2025). GaussianSlicer: Efficient Surface Reconstruction from Cross-sectional Slices with Gaussian Splatting. In ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 1–5). IEEE. ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). https://doi.org/10.1109/icassp49660.2025.10890834

[8] J. Li, T. B. Culver, P. P. Persaud, and J. M. Hathaway, "Developing nitrogen removal models for stormwater bioretention systems," Water Research, vol. 243, p. 120381, 2023.

[9] J. Li and T. B. Culver, "Review of process-based nitrogen model for agricultural fields with implications for nitrogen simulations in stormwater BMPs," Environmental Modelling & Software, vol. 151, p. 105363, 2022.

[10] Li, T. B. Culver, C. R. Burgis, W. Zhang, and J. A. Smith, "Validating Nitrogen Removal Models with Field Bioretention Data," Journal of Environmental Engineering, vol. 150, no. 8, p. 04024037, 2024.

[11] Li, "Nitrogen Removal Models for Stormwater Bioretention Systems," Ph.D. dissertation, University of Virginia, 2023.

[12] Liang, J., Wang, Z., Ma, Z., Li, J., Zhang, Z., Wu, X., & Wang, B. (2024). Online training of large language models: Learn while chatting. arXiv preprint arXiv:2403.04790.

[13] Wang, Z., Su, J., Zhou, M., Zeng, H., Jia, M., Lv, X., ... & Zhang, D. (2025). SheetBrain: A Neuro-Symbolic Agent for Accurate Reasoning over Complex and Large Spreadsheets. arXiv preprint arXiv:2510.19247.

[14] Ge, W., Wang, Z., Wang, P., Liang, J., Cai, Z. G., Mai, Z., & Wang, B. (2024). Towards Gamifying Interactive Language Learning using Large Language Models for Children.

[15] Wang, Z., Zhang, Q., Liu, T., & Li, C. (2024). Analyzing Financial News Sentiment with NLP to Forecast Market Trends. International Journal of Engineering and Management Research, 14(5), 6-11.

[16] Rao, Jiarui, et al. "Optimizing Stock Market Return Forecasts with Uncertainty Sentiment: Leveraging LLM-based Insights." Proceedings of the 2024 5th International Conference on Big Data Economy and Information Management. 2024.

[17] Rao, Jiarui, and Jionghao Lin. "Ramo: Retrieval-augmented generation for enhancing moocs recommendations." arXiv preprint arXiv:2407.04925 (2024).

[18] Rao, Jiarui, Qian Zhang, and Xinqiu Liu. "Applications Analyzing E-commerce Reviews with Large Language Models (LLMs): A Methodological Exploration and Application Insight." Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023 7.01 (2024): 207-212.
[19] Zhang, Qian, et al. "Sea MNF vs. LDA: Unveiling the power of short text mining in financial markets." International Journal of Engineering and Management Research 14.5 (2024): 76-82.
[20] Rao, Jiarui, et al. "Integrating Textual Analytics with Time Series Forecasting Models: Enhancing Predictive Accuracy in Global Energy and Commodity Markets." Innovations in Applied Engineering and Technology (2023): 1-7.
[21] Zhang, Qian, and Jiarui Rao. "Enhancing Financial Forecasting Models with Textual Analysis: A Comparative Study of Decomposition Techniques and Sentiment-Driven Predictions." Innovations in Applied Engineering and Technology (2022): 1-6.
[22] Lin, Jionghao, et al. "Automatic large language models creation of interactive learning lessons." European Conference on Technology Enhanced Learning. Cham: Springer Nature Switzerland, 2025.
[23] Peng, Jingyang, et al. "Automated bias assessment in ai-generated educational content using ceat framework." International Conference on Artificial Intelligence in Education. Cham: Springer Nature Switzerland, 2025.
[24] Rao, Jiarui, and Qian Zhang. "Deconstructing Digital Discourse: A Deep Dive into Distinguishing LLM-Powered Chatbots from Human Language." Journal of Theory and Practice in Education and Innovation 2.2 (2025): 18-25.